

Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations

Benjamin A. Burton

January 16, 2011

Abstract

Enumerating all 3-manifold triangulations of a given size is a difficult but increasingly important problem in computational topology. A key difficulty for enumeration algorithms is that most combinatorial triangulations must be discarded because they do not represent topological 3-manifolds. In this paper we show how to preempt bad triangulations by detecting genus in partially-constructed vertex links, allowing us to prune the enumeration tree substantially.

The key idea is to manipulate the boundary edges surrounding partial vertex links using expected logarithmic time operations. Practical testing shows the resulting enumeration algorithm to be significantly faster, with up to $249\times$ speed-ups even for small problems where comparisons are feasible. We also discuss parallelisation, and describe new data sets that have been obtained using high-performance computing facilities.

ACM classes G.2.1; G.4; I.1.2

Keywords Computational topology, 3-manifolds, triangulations, census algorithm, combinatorial enumeration

1 Introduction

In computational geometry and topology, triangulations are natural and ubiquitous data structures for representing topological spaces. Here we focus on triangulations in 3-manifold topology, an important branch of topology in which many key problems are theoretically decidable but extremely difficult for practical computation [1, 11].

A *census* of 3-manifold triangulations is a list of all triangulations that satisfy some given set of properties. A typical census fixes the number of tetrahedra (the *size* of the triangulation), and enumerates all triangulations up to *isomorphism* (a relabelling of the tetrahedra and their vertices).

Censuses of this type first appeared in the late 1980s [14, 19].¹ These were censuses of *minimal triangulations*, which represent a given 3-manifold using the fewest possible tetrahedra. Such censuses have brought about new insights into the combinatorics of minimal triangulations [9, 18, 21] and the complexities of 3-manifolds [17, 20], and have proven useful for computation and experimentation [4, 10].

A more recent development has been censuses of *all* possible 3-manifold triangulations of a given size, including non-minimal triangulations [11]. These have yielded surprising experimental

¹Some authors, following Matveev [19], work in the setting of *special spines* which are dual to triangulations.

insights into algorithmic complexity problems and random 3-manifold triangulations [11, 12], both topics which remain extremely difficult to handle theoretically.

The limits of censuses in the literature are fairly small. For closed \mathbb{P}^2 -irreducible 3-manifolds, all minimal triangulations have been enumerated for only ≤ 10 tetrahedra [8]; for the orientable case only, the list of manifolds (but not triangulations) is known for ≤ 12 tetrahedra [23]. A full list of all closed 3-manifold triangulations (non-minimal included) is known for just ≤ 9 tetrahedra [11]. These small limits are unavoidable because censuses grow exponentially—and sometimes super-exponentially [12]—in size.

Nevertheless, in theory it should be possible to extend these results substantially. Census algorithms are typically based on a recursive search through all possible *combinatorial triangulations*—that is, methods of gluing together faces of tetrahedra in pairs. However, as the number of tetrahedra grows large, almost all combinatorial triangulations are *not* 3-manifold triangulations [13]. The problem is that *vertex links*—boundaries of small neighbourhoods of the vertices of a triangulation—are generally not spheres or discs as they should be, but instead higher-genus surfaces.

To illustrate: for $n = 9$ tetrahedra, if we simply glue together all $4n$ tetrahedron faces in pairs, a very rough estimate (described in the appendix) gives at least 6.44×10^{12} connected *combinatorial* triangulations up to isomorphism. However, just 139 103 032 are *3-manifold* triangulations [11], and a mere 3 338 are minimal triangulations of closed \mathbb{P}^2 -irreducible 3-manifolds [8].

It is clear then that large branches of the combinatorial search tree can be avoided, if one could only identify *which* branches these are. The current challenge for enumeration algorithms is to find new and easily-testable conditions under which such branches can be pruned.

For censuses of minimal triangulations, several such conditions are known: examples include the absence of low-degree edges [6, 14, 21], or of “bad subgraphs” in the underlying 4-valent face pairing graph [8, 16]. Nevertheless, seeing how few minimal triangulations are found in practice, there are likely many more conditions yet to be found.

It is critical that such conditions can be tested quickly, since these tests are run on a continual basis as the search progresses and backtracks. The paper [8] introduces a modified union-find framework through which several minimality tests can be performed in $O(\log n)$ time. Importantly, this framework can also be used with censuses of *all* 3-manifold triangulations (non-minimal included), where it is used to test that: (i) partially-constructed vertex links are orientable, and (ii) fully-constructed edges are not identified with themselves in reverse. Although tests (i) and (ii) are powerful when enumerating *non-orientable* triangulations, they do not help for *orientable* triangulations because they are already enforced by the overall enumeration algorithm.

The main contributions of this paper are:

- We add the following condition to the suite of tests used during enumeration: *all partially-constructed vertex links must be punctured spheres* (not punctured higher-genus surfaces).

Although this condition is straightforward, it has traditionally required $O(n)$ operations to test, making it impractical for frequent use in the enumeration algorithm. In Section 3 we show how to test this condition *incrementally*, using only expected logarithmic-time operations at each stage of the combinatorial search.

This condition is extremely powerful in both the orientable and non-orientable cases, and our new incremental test makes it practical for real use. Performance testing in Section 4

shows speed-ups of up to $249\times$ even for small enumeration problems where experimental comparisons are feasible ($n \leq 7$).

- In Section 6 we use this to obtain new census data, including all closed 3-manifold triangulations of size $n \leq 10$ (non-minimal included, improving the previous limit of $n \leq 9$), and all minimal triangulations of closed \mathbb{P}^2 -irreducible 3-manifolds of size $n \leq 11$ (improving the previous limit of $n \leq 10$). High-performance computing and distributed algorithms play a key role, as outlined in Section 5. All censuses cover both orientable and non-orientable triangulations.

This new census data is already proving useful in ongoing projects, such as studying the structure of minimal triangulations, and mapping out average-case and generic complexities for difficult topological decision problems.

It should be noted that avoiding isomorphisms—often a significant difficulty in combinatorial enumeration—is not a problem here. See the full version of this paper for details.

Looking forward, the techniques of this paper can also be applied to the enumeration of *4-manifold* triangulations. In this higher-dimensional setting, our techniques can be applied to edge links rather than vertex links. Vertices on the other hand become more difficult to handle: each vertex link must be a 3-sphere, and 3-sphere recognition remains a difficult algorithmic problem [12, 20]. Here research into algebraic techniques may yield new heuristics to further prune the search tree.

Throughout this paper we restrict our attention to closed 3-manifolds, although all of the results presented here extend easily to manifolds with boundary.

All algorithms described in this paper can be downloaded as part of *Regina* [5, 7], an open-source software package for the algebraic and combinatorial manipulation of 3-manifolds and their triangulations.

2 Preliminaries

Consider a collection of n tetrahedra (these are abstract objects, and need not be embedded in some \mathbb{R}^d). A *combinatorial triangulation* of size n is obtained by affinely identifying (or “gluing together”) the $4n$ tetrahedron faces in pairs.² Specifically, it consists of the following data:

- a partition of the $4n$ tetrahedron faces into $2n$ pairs, indicating which faces are to be identified;
- $2n$ permutations of three elements, indicating which of the six possible rotations or reflections is to be used for each identification.

For instance, consider the following example with $n = 3$ tetrahedra. The tetrahedra are labelled A, B, C , and the four vertices of each tetrahedron are labelled $0, 1, 2, 3$.

Tetrahedron	Face 012	Face 013	Face 023	Face 123
A	$C : 013$	$B : 012$	$A : 312$	$A : 230$
B	$A : 013$	$C : 120$	$C : 231$	$C : 302$
C	$B : 301$	$A : 012$	$B : 231$	$B : 302$

²A combinatorial triangulation need not be a simplicial complex, and need not represent a topological 3-manifold. The word “combinatorial” indicates that we are only interested in face identifications, with no topological requirements.

The top-left cell of this table indicates that face 012 of tetrahedron A is identified with face 013 of tetrahedron C , using the rotation or reflection that maps vertices 0, 1, 2 of tetrahedron A to vertices 0, 1, 3 of tetrahedron C respectively. For convenience, the same identification is also shown from the other direction in the second cell of the bottom row.

As a consequence of these face identifications, we find that several tetrahedron edges become identified together; each such equivalence class is called an *edge of the triangulation*. Likewise, each equivalence class of identified vertices is called a *vertex of the triangulation*. The triangulation illustrated above has three edges and just one vertex.

The *face pairing graph* of a combinatorial triangulation is the 4-valent multigraph whose nodes represent tetrahedra and whose edges represent face identifications. The face pairing graph for the example above is shown in Figure 1. A combinatorial triangulation is called *connected* if and only if its face pairing graph is connected.




Figure 1: An example face pairing graph

The *vertex links* of a triangulation are obtained as follows. In each tetrahedron we place four triangles surrounding the four vertices, as shown in Figure 2. We then glue together the edges of these triangles in a manner consistent with the face identifications of the surrounding tetrahedra, as illustrated in Figure 3.




Figure 2: The triangles that form the vertex links




Figure 3: Joining vertex linking triangles along their edges in adjacent tetrahedra

The result is a collection of triangulated closed surfaces, one surrounding each vertex of the triangulation. The surface surrounding vertex V is referred to as the *link of V* . Topologically, this represents the boundary of a small neighbourhood of V in the triangulation.

A *3-manifold triangulation* is a combinatorial triangulation that, when viewed as a topological space, represents a 3-manifold. Equivalently, a 3-manifold triangulation is a combinatorial triangulation in which:

- (i) each vertex link is a topological sphere;

(ii) no tetrahedron edge is identified with itself in reverse as a result of the face identifications.³

The earlier example is *not* a 3-manifold triangulation, since the link of the (unique) vertex is a torus, not a sphere. For many 3-manifolds M , the size of a minimal triangulation of M corresponds to the Matveev complexity of M [17, 20].

A *partial triangulation* is a combinatorial triangulation in which we identify only $2k$ of the $4n$ tetrahedron faces in pairs, for some $0 \leq 2k \leq 4n$. We define vertices, edges and vertex links as before, noting that vertex links might now be surfaces with boundary (not closed surfaces).

A typical enumeration algorithm works as follows [8, 14]:

Algorithm 1. *Suppose we wish to enumerate all connected 3-manifold triangulations of size n satisfying some set of properties P . The main steps are:*

1. *Enumerate all possible face pairing graphs (i.e., all connected 4-valent multigraphs on n nodes).*
2. *For each graph, recursively try all 6^{2n} possible rotations and reflections for identifying the corresponding tetrahedron faces. Each “partial selection” of rotations and reflections gives a partial triangulation, and recursion and backtracking correspond to gluing and ungluing tetrahedron faces in these partial triangulations.*
3. *Whenever we have a partial triangulation, run a series of tests that can identify situations where, no matter how we glue the remaining faces together, we can never obtain a 3-manifold triangulation satisfying the properties in P . If this is the case, prune the current branch of the search tree and backtrack immediately.*
4. *Whenever we have a complete selection of 6^{2n} rotations and reflections, test whether (i) the corresponding combinatorial triangulation is in fact a 3-manifold triangulation, and (ii) whether this 3-manifold triangulation satisfies the required properties in P .*

There is also the problem of avoiding isomorphisms. This is computationally cheap if the recursion is ordered carefully; for details, see the full version of this paper.

In practice, step 1 is negligible—almost all of the computational work is in the recursive search (steps 2–4). The tests in step 3 are critical: they must be extremely fast, since they are run at every stage of the recursive search. Moreover, if chosen carefully, these tests can prune vast sections of the search tree and speed up the enumeration substantially.

A useful observation is that some graphs can be eliminated immediately after step 1. See [6, 8, 16] for algorithms that incorporate such techniques.

3 Tracking Vertex Links

In this paper we add the following test to step 3 of the enumeration algorithm:

Test 2. *Whenever we have a partial triangulation \mathcal{T} , test whether all vertex links are spheres with zero or more punctures. If not, prune the current branch of the search tree and backtrack immediately.*

Theoretically, it is simple to show that this test works:

³An equivalent condition to (ii) is that we can direct the edges of every tetrahedron in a manner consistent with the face identifications.

Lemma 3. *If \mathcal{T} is a partial triangulation and the link of some vertex V is not a sphere with zero or more punctures, then there is no way to glue together the remaining faces of \mathcal{T} to obtain a 3-manifold triangulation.*

Proof. Suppose we *can* glue the remaining faces together to form a 3-manifold triangulation \mathcal{T}' . Let L and L' be the links of V in \mathcal{T} and \mathcal{T}' respectively; since \mathcal{T}' is a 3-manifold triangulation, L' must be a topological sphere.

This link L' is obtained from L by attaching zero or more additional triangles. Therefore L is an embedded subsurface of the sphere, and so L must be a sphere with zero or more punctures. \square

The test itself is straightforward; the difficulty lies in performing it *quickly*. A fast implementation is crucial, since it will be called repeatedly throughout the recursive search.

The key idea is to track the 1-dimensional *boundary curves* of the vertex links, which are formed from cycles of edges belonging to vertex-linking triangles. As we glue tetrahedron faces together, we repeatedly split and splice these boundary cycles. To verify Test 2, we must track which triangles belong to the same vertex links and which edges belong to the same boundary cycles, which we can do in expected logarithmic time using union-find and skip lists respectively.

In the sections below, we describe what additional data needs to be stored (Section 3.1), how to manipulate and use this data (Section 3.2), and how skip lists can ensure a small time complexity (Section 3.3).

3.1 Data structures

In a partial triangulation with n tetrahedra, there are $4n$ *vertex linking triangles* that together form the vertex links (four such triangles are shown in Figure 2). These $4n$ triangles are surrounded by a total of $12n$ *vertex linking edges*.

Each time we glue together two tetrahedron faces, we consequently glue together three pairs of vertex linking edges, as shown in Figure 3. This gradually combines the triangles into a collection of larger triangulated surfaces, as illustrated in Figure 4. The boundary curves of these surfaces are drawn in bold in Figure 4; these are formed from the vertex linking edges that have not yet been paired together.




Figure 4: Vertex linking surfaces after gluing several tetrahedron faces

To support Test 2, we store all $12n$ vertex linking edges in a series of cyclic list structures that describe these boundary curves. To simplify the discussion, we begin with a naïve implementation based on doubly-linked lists. However, this leaves us with an $O(n)$ operation to perform, as seen in Section 3.2. To run all operations in expected logarithmic time we use skip lists [25], which we outline in Section 3.3.

We treat the vertex linking edges as *directed edges* (i.e., arrows), with directions chosen arbitrarily at the beginning of the enumeration algorithm. For each vertex linking edge e , we store the following data:

- If e is part of a boundary curve, we store the two edges adjacent to e along this boundary curve, as well as two booleans that tell us whether these adjacent edges point in the same or opposite directions.
- If e is not part of a boundary curve (i.e., it has been glued to some other vertex linking edge and is now internal to a vertex linking surface), we store a snapshot of the above data from the last time that e *was* part of a boundary curve.

To summarise: edges on the boundary curves are stored in a series of doubly-linked lists, and internal edges remember where they *were* in these lists right before they were glued to their current partner.

3.2 Recursion, backtracking and testing

Recall that each time we glue two tetrahedron faces together, we must glue together *three* pairs of vertex linking edges. Each of these edge gluings changes the vertex linking surfaces, and so we process each edge gluing individually.

There are three key operations that we must perform in relation to edge gluings:

- (i) gluing two vertex linking edges together (when we step forward in the recursion);
- (ii) ungluing two vertex linking edges (when we backtrack);
- (iii) verifying Test 2 after gluing two vertex linking edges together (i.e., verifying that all vertex links are spheres with zero or more punctures).

We now present the details of each operation in turn. Throughout this discussion we assume that edges are glued together so that all vertex links are *orientable* surfaces; the paper [8] describes an efficient framework for detecting non-orientable vertex links as soon as they arise.

Recursion: gluing edges together

Suppose we wish to glue together edges x and y , as illustrated in Figure 5. Only local modifications are required: edges p and r become adjacent and must now to link to each other (instead of to x and y); likewise, edges q and s must be adjusted to link to each other. Note that this gluing introduces a change of direction where p and r meet (and likewise for q and s), so as we adjust the direction-related booleans we must perform an extra negation on each side.




Figure 5: Gluing two vertex linking edges together

We make no changes to the data stored for edges x and y , since these two edges are now internal and their associated data now represents a snapshot from the last time that they were boundary edges (as required by Section 3.1).

All of these local modifications can be performed in $O(1)$ time. The associated list operations are deletion, splitting and splicing; this becomes important when we move to skip lists in Section 3.3.

It is important to remember the special case in which edges x and y are adjacent in the same boundary cycle. Here the local modifications are slightly different (there are only two or possibly zero nearby edges to update instead of four), but these modifications remain $O(1)$ time.

Backtracking: ungluing edges

As with gluing, ungluing a pair of vertex linking edges is a simple matter of local modifications. Here the backtracking context is important: it is essential that we unglue edges in the reverse order to that in which they were glued.

Suppose we are ungluing edges x and y as depicted in Figure 5. The snapshot data stored with edges x and y shows that they *were* adjacent to edges p, q, r and s immediately before this gluing was made (and therefore immediately *after* the ungluing that we are now performing).

This snapshot data therefore gives us access to edges p, q, r and s : now we simply adjust p and q to link to x (instead of r and s), and likewise we adjust r and s to link to y . No modifications to edges x and y are required.

Again we must adjust the direction-related booleans carefully, and we must cater for the case in which edges x and y were adjacent immediately before the gluing was made.

As before, all local modifications can be performed in $O(1)$ time. For the skip list discussion in Section 3.3, the associated list operations are splitting, splicing and insertion.

Testing: verifying that links are punctured spheres

Each time we glue two vertex linking edges together we must ensure that every vertex link is a sphere with zero or more punctures (Test 2). We test this *incrementally*: we assume this is true *before* we glue these edges together, and we verify that our new gluing does not introduce any unwanted genus to the vertex links.

Our incremental test is based on the following two observations:

Lemma 4. *Let S and S' be distinct triangulated spheres with punctures, and let x and y be boundary edges from S and S' respectively. If we glue x and y together, the resulting surface is again a sphere with punctures.*

Lemma 5. *Let S be a triangulated sphere with punctures, and let x and y be distinct boundary edges of S . If we glue x and y together in an orientation-preserving manner, the resulting surface is a sphere with punctures if and only if x and y belong to the same boundary cycle of S .*

The proofs of Lemmata 4 and 5 are simple, and we do not give further details here. Figure 6(a) illustrates the scenario of Lemma 4 with two distinct punctured spheres, and Figures 6(b)–6(d) show different scenarios from Lemma 5 in which we join two boundary edges from the same punctured sphere.

In particular, Figure 6(b) shows the case in which x and y belong to different boundary cycles; here we observe that the resulting surface is a twice-punctured torus. Figures 6(c) and 6(d) show cases with x and y on the same boundary cycle; in 6(d), x and y are adjacent along the boundary.

Note that Lemmata 4 and 5 hold even with very short boundaries (for instance, one-edge boundaries consisting of x or y alone).




Figure 6: Different ways of gluing boundary edges together

It is now clear how to incrementally verify Test 2 when we glue together vertex linking edges x and y :

1. Test whether the vertex linking triangles containing x and y belong to the same connected vertex linking surface. If not, the test passes. Otherwise:
2. Test whether x and y belong to the same doubly-linked list of boundary edges (i.e., the same boundary cycle). If so, the test passes. If not, the test fails.

Here we implicitly assume that all gluings are orientation-preserving, as noted at the beginning of Section 3.2.

Step 1 can be performed in $O(\log n)$ time using the modified union-find structure outlined in [8]. The original purpose of this structure was to enforce orientability in vertex linking surfaces, and one of the operations it provides is an $O(\log n)$ test for whether two vertex linking triangles belong to the same connected vertex linking surface. This modified union-find supports backtracking; see [8] for further details.

Step 2 is more difficult: a typical implementation might involve walking through the doubly-linked list containing x until we either find y or verify that y is not present, which takes $O(n)$ time to complete. Union-find cannot help us, because of our repeated splitting and splicing of boundary cycles. In the following section we show how to reduce this $O(n)$ running time to expected $O(\log n)$ by extending our doubly-linked lists to become *skip lists*.

It should be noted that Step 2 can in fact be carried out in $O(b)$ time, where b is the number of boundary edges on all vertex linking surfaces. Although $b \in O(n)$ in general, for some face pairing graphs b can be far smaller. For instance, when the face pairing graph is a double-ended chain [6], we can arrange the recursive search so that $b \in O(1)$. See the full version of this paper for details.

3.3 Skip lists and time complexity

From the discussion in Section 3.2, we see that with our naïve doubly-linked list implementation, the three key operations of gluing edges, ungluing edges and verifying Test 2 have $O(1)$, $O(1)$ and $O(n)$ running times respectively. The bottleneck is the $O(n)$ test for whether two vertex linking edges belong to the same doubly-linked list (i.e., the same boundary cycle).

We can improve our situation by extending our doubly-linked list to a *skip list* [25]. Skip lists are essentially linked lists with additional layers of “express pointers” that allow us to move quickly through the list instead of stepping forward one element at a time. Figure 7 shows a typical skip list structure.




Figure 7: The internal layout of a skip list

The list operations used in Section 3.2 for gluing and ungluing edges are deletion, insertion, splitting and splicing; all of these can be performed on a skip list in expected $O(\log n)$ time [24, 25]. Importantly, it also takes expected $O(\log n)$ time to search forward to the last element of a skip list.⁴ We can therefore test whether vertex linking edges x and y belong to the same list by searching forward to the end of each list and testing whether the final elements are the same.

It follows that, with a skip list implementation, all three key operations of gluing edges, ungluing edges and verifying Test 2 run in expected $O(\log n)$ time. Therefore:

Theorem 6. *It is possible to implement Test 2 by performing expected $O(\log n)$ operations at every stage of the recursive search.*

Note that we must include reverse links at the lowest layer of the skip list (effectively maintaining the original doubly-linked list), since both forward and backward links are required for the gluing and ungluing operations (see Section 3.2). Full details of the skip list implementation can be found in the full version of this paper.

4 Performance

Here we measure the performance of our new algorithm experimentally. Specifically, we compare two enumeration algorithms: the *old algorithm*, which includes all of the optimisations described in [8] (including the union-find framework for ensuring that vertex links remain orientable), and the *new algorithm*, which enhances the old algorithm with the new tests described in Section 3.⁵

⁴Although our lists are cyclic, we can always define an arbitrary endpoint.

⁵We compare the new algorithm against [8] because this allows us to isolate our new techniques, and because the source code and implementation details for alternative algorithms [16, 20] are not readily available.

We run our performance tests by enumerating censuses of all 3-manifold triangulations of size $n \leq 7$. We choose this type of census because a census of minimal triangulations requires significant manual post-processing [8], and because a census of all triangulations is significantly larger and therefore a stronger “stress test”. We restrict our tests to $n \leq 7$ because for larger n the old algorithm becomes too slow to run time trials on a single CPU.

Census parameters	Triangulations	Old (h:m:s)	New (m:s)	Speed-up
$n = 5$, orientable	4 807	0:59	0:03	20×
$n = 5$, non-orientable	377	1:09	0:06	12×
$n = 6$, orientable	52 946	1:11:57	1:03	69×
$n = 6$, non-orientable	4 807	1:23:30	2:05	40×
$n = 7$, orientable	658 474	92:23:39	22:16	249×
$n = 7$, non-orientable	64 291	103:24:51	48:27	128×

Table 1: Running times for old and new algorithms

Table 1 shows the results of our time trials, split into censuses of orientable and non-orientable triangulations. For $n \leq 4$ both algorithms run in 1 second or less. All trials were carried out on a single 2.93 GHz Intel Xeon X5570 CPU.

The results are extremely pleasing: for $n = 7$ we see a speed-up of 128× in the non-orientable case and 249× in the orientable case (from almost four days of running time down to just 22 minutes). Moreover, the speed-up factors appear to grow exponentially with n . All of this suggests that our new algorithm can indeed make a concrete difference as to how large a census we can feasibly build.

It is worth noting that speed-ups are consistently better for the orientable case. This may be because the union-find framework introduced in [8] is most effective for non-orientable enumeration (as noted in Section 1), and so the orientable case has more room for gain. Nevertheless, it is pleasing to see that the new algorithm gives substantial improvements for both the orientable and non-orientable cases.

5 Parallelisation

As n increases, the output size for a typical census grows exponentially in n , and sometimes super-exponentially—for instance, the growth rate of a census of all 3-manifold triangulations is known to be $\exp(\Theta(n \log n))$ [12]. It is therefore critical that enumeration algorithms be parallelised if we are to make significant progress in obtaining new census data.

Like many combinatorial searches, the enumeration of triangulations is an embarrassingly parallel problem: different branches of the search tree can be processed independently, making the problem well-suited for clusters and server farms. Avoiding isomorphisms causes some minor complications, which we discuss in the full version of this paper.

The main obstacle is that, because of the various pruning techniques (as described in Sections 2 and 3), it is very difficult to estimate in advance how long each branch of the search tree will take to process. Experience shows that there can be orders-of-magnitude differences in running time between subsearches at the same depth in the tree.

For this reason, parallelisation must use a controller/slave model in which a controller process repeatedly hands small pieces of the search space to the next available slave, as opposed

to a simple subdivision in which each process handles a fixed portion of the search space. This means that some inter-process communication is required.

For each subsearch, the controller must send $O(n)$ data to the slave: this includes the face pairing graph, the partial triangulation, and the data associated with the vertex linking edges and triangles as described in Section 3. The output for each subsearch can be super-exponentially large, and so it is preferable for slaves to write this data directly to disk (as opposed to communicating it back to the controller). Collating the output from different slaves is a simple task that can be performed after the enumeration has finished.

The enumeration code in *Regina* implements such a model using MPI, and runs successfully on hundreds of simultaneous CPUs with a roughly proportional speed-up in wall time.

6 Census Data

The new algorithms in this paper have been implemented and run in parallel using high-performance computing facilities to obtain new census data that exceeds the best known limits in the literature. This includes (i) a census of all closed 3-manifold triangulations of size $n \leq 10$, and (ii) a census of all minimal triangulations of closed \mathbb{P}^2 -irreducible 3-manifolds of size $n \leq 11$.

6.1 All closed 3-manifold triangulations

The first reported census of all closed 3-manifold triangulations appears in [11] for $n \leq 9$, and has been used to study algorithmic complexity and random triangulations [11, 12]. Here we extend this census to $n \leq 10$ with a total of over 2 billion triangulations:

Theorem 7. *There are precisely 2 196 546 921 closed 3-manifold triangulations that can be constructed from ≤ 10 tetrahedra, as summarised by Table 2.*

Size (n)	Orientable	Non-orientable	Total
1	4	—	4
2	16	1	17
3	76	5	81
4	532	45	577
5	4 807	377	5 184
6	52 946	4 807	57 753
7	658 474	64 291	722 765
8	8 802 955	984 554	9 787 509
9	123 603 770	15 499 262	139 103 032
10	1 792 348 876	254 521 123	2 046 869 999
Total	1 925 472 456	271 074 465	2 196 546 921

Table 2: All closed 3-manifold triangulations

The total CPU time required to enumerate the 10-tetrahedron census was $\simeq 2.4$ years, divided amongst 192 distinct 2.93 GHz Intel Xeon X5570 CPUs.

The paper [11] makes two conjectures regarding the worst-case and average number of vertex normal surfaces for a closed 3-manifold triangulation of size n . Details and definitions can be found in [11]; in summary, these conjectures are:

Conjecture 1. *For all positive $n \neq 1, 2, 3, 5$, a tight upper bound on the number of vertex normal surfaces in a closed 3-manifold triangulation of size n is:*

$$\begin{array}{ll} 17^k + k & \text{if } n = 4k; \\ 581 \cdot 17^{k-2} + k + 1 & \text{if } n = 4k + 1; \\ 69 \cdot 17^{k-1} + k & \text{if } n = 4k + 2; \\ 141 \cdot 17^{k-1} + k + 2 & \text{if } n = 4k + 3, \end{array}$$

and so this upper bound grows asymptotically as $\Theta(17^{n/4})$.

Conjecture 2. *If $\bar{\sigma}_n$ represents the average number of vertex normal surfaces amongst all closed 3-manifold triangulations (up to isomorphism), then $\bar{\sigma}_n < \bar{\sigma}_{n-1} + \bar{\sigma}_{n-2}$ for all $n \geq 3$, and so $\bar{\sigma}_n \in O([\frac{1+\sqrt{5}}{2}]^n)$.*

These conjectures were originally based on the census data for $n \leq 9$. With our new census we can now verify these conjectures at the 10-tetrahedron level:

Theorem 8. *Conjectures 1 and 2 are true for all $n \leq 10$.*

This census contains over 63 GB of data, and so the data files have not been posted online. Readers who wish to work with this data are welcome to contact the author for a copy.

6.2 Closed \mathbb{P}^2 -irreducible 3-manifolds

A 3-manifold is \mathbb{P}^2 -irreducible if every sphere bounds a ball and there are no embedded two-sided projective planes. Censuses of closed \mathbb{P}^2 -irreducible 3-manifolds and their minimal triangulations have a long history [2, 3, 8, 9, 15, 16, 22, 23, 19]. The largest reported census of all minimal triangulations of these manifolds reaches $n \leq 10$ [8]. If we enumerate *manifolds* but not their triangulations, the census reaches $n \leq 12$ in the orientable case [23] but remains at $n \leq 10$ in the non-orientable case.

Here we extend this census of minimal triangulations of closed \mathbb{P}^2 -irreducible 3-manifolds to $n \leq 11$. As a result, we also extend the census of underlying manifolds to $n \leq 11$ in the non-orientable case, and in the orientable case we confirm that the number of manifolds matches Matveev's census [23].

Theorem 9. *There are precisely 13 765 closed \mathbb{P}^2 -irreducible 3-manifolds that can be constructed from ≤ 11 tetrahedra. These have a combined total of 55 488 minimal triangulations, as summarised by Table 3.*

The paper [9] raises conjectures for certain classes of non-orientable 3-manifolds regarding the combinatorial structure of every minimal triangulation. Again we refer to the source [9] for details and definitions; in summary:

Conjecture 3. *Every minimal triangulation of a non-flat non-orientable torus bundle over the circle is a layered torus bundle.*

Conjecture 4. *Every minimal triangulation of a non-flat non-orientable Seifert fibred space over $\mathbb{R}P^2$ or \bar{D} with two exceptional fibres is either a plugged thin I -bundle or a plugged thick I -bundle.*

Size (n)	Minimal triangulations		Distinct 3-manifolds	
	Orientable	Non-orient.	Orientable	Non-orient.
1	4	—	3	—
2	9	—	6	—
3	7	—	7	—
4	15	—	14	—
5	40	—	31	—
6	115	24	74	5
7	309	17	175	3
8	945	59	436	10
9	3 031	307	1 154	33
10	10 244	983	3 078	85
11	36 097	3 282	8 421	230
Total	50 816	4 672	13 399	366

Table 3: All minimal triangulations of closed \mathbb{P}^2 -irreducible 3-manifolds

Layered torus bundles and plugged thin and thick I -bundles are families of triangulations with well-defined combinatorial structures. The original conjectures were based on census data for $n \leq 8$, and in [8] they are shown to hold for all $n \leq 10$. With our new census data we are now able to validate these conjectures at the 11-tetrahedron level:

Theorem 10. *Conjectures 3 and 4 are true for all minimal triangulations of size $n \leq 11$.*

Data files for this census, including the 3-manifolds and all of their minimal triangulations, can be downloaded from the *Regina* website [5].

Acknowledgments

Computational resources used in this work were provided by the Queensland Cyber Infrastructure Foundation and the Victorian Partnership for Advanced Computing.

References

- [1] Ian Agol, Joel Hass, and William Thurston, *3-manifold knot genus is NP-complete*, STOC '02: Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing, ACM Press, 2002, pp. 761–766.
- [2] Gennaro Amendola and Bruno Martelli, *Non-orientable 3-manifolds of small complexity*, Topology Appl. **133** (2003), no. 2, 157–178.
- [3] ———, *Non-orientable 3-manifolds of complexity up to 7*, Topology Appl. **150** (2005), no. 1-3, 179–195.
- [4] Ryan Budney, *Embeddings of 3-manifolds in S^4 from the point of view of the 11-tetrahedron census*, Preprint, [arXiv:0810.2346](https://arxiv.org/abs/0810.2346), October 2008.
- [5] Benjamin A. Burton, *Regina: Normal surface and 3-manifold topology software*, <http://regina.sourceforge.net/>, 1999–2010.
- [6] ———, *Face pairing graphs and 3-manifold enumeration*, J. Knot Theory Ramifications **13** (2004), no. 8, 1057–1101.

- [7] ———, *Introducing Regina, the 3-manifold topology software*, Experiment. Math. **13** (2004), no. 3, 267–272.
- [8] ———, *Enumeration of non-orientable 3-manifolds using face-pairing graphs and union-find*, Discrete Comput. Geom. **38** (2007), no. 3, 527–571.
- [9] ———, *Observations from the 8-tetrahedron nonorientable census*, Experiment. Math. **16** (2007), no. 2, 129–144.
- [10] ———, *Converting between quadrilateral and standard solution sets in normal surface theory*, Algebr. Geom. Topol. **9** (2009), no. 4, 2121–2174.
- [11] ———, *The complexity of the normal surface solution space*, SCG '10: Proceedings of the Twenty-Sixth Annual Symposium on Computational Geometry, ACM, 2010, pp. 201–209.
- [12] ———, *The Pachner graph and the simplification of 3-sphere triangulations*, To appear in SCG '11: Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry, [arXiv:1011.4169](https://arxiv.org/abs/1011.4169), November 2010.
- [13] Nathan M. Dunfield and William P. Thurston, *Finite covers of random 3-manifolds*, Invent. Math. **166** (2006), no. 3, 457–521.
- [14] Martin V. Hildebrand and Jeffrey R. Weeks, *A computer generated census of cusped hyperbolic 3-manifolds*, Computers and Mathematics (Cambridge, MA, 1989), Springer, New York, 1989, pp. 53–59.
- [15] Bruno Martelli, *Complexity of 3-manifolds*, Spaces of Kleinian Groups, London Math. Soc. Lecture Note Ser., vol. 329, Cambridge Univ. Press, Cambridge, 2006, pp. 91–120.
- [16] Bruno Martelli and Carlo Petronio, *Three-manifolds having complexity at most 9*, Experiment. Math. **10** (2001), no. 2, 207–236.
- [17] ———, *A new decomposition theorem for 3-manifolds*, Illinois J. Math. **46** (2002), 755–780.
- [18] ———, *Complexity of geometric three-manifolds*, Geom. Dedicata **108** (2004), no. 1, 15–69.
- [19] S. V. Matveev and A. T. Fomenko, *Constant energy surfaces of Hamiltonian systems, enumeration of three-dimensional manifolds in increasing order of complexity, and computation of volumes of closed hyperbolic manifolds*, Russian Math. Surveys **43** (1988), no. 1, 3–24.
- [20] Sergei Matveev, *Algorithmic topology and classification of 3-manifolds*, Algorithms and Computation in Mathematics, no. 9, Springer, Berlin, 2003.
- [21] Sergei V. Matveev, *Tables of 3-manifolds up to complexity 6*, Max-Planck-Institut für Mathematik Preprint Series (1998), no. 67, available from <http://www.mpim-bonn.mpg.de/html/preprints/preprints.html>.
- [22] ———, *Recognition and tabulation of three-dimensional manifolds*, Dokl. Akad. Nauk **400** (2005), no. 1, 26–28.
- [23] ———, *Tabulation of three-dimensional manifolds*, Russian Math. Surveys **60** (2005), no. 4, 673–698.
- [24] William Pugh, *A skip list cookbook*, Report UMIACS-TR-89-72.1, Univ. of Maryland Institute for Advanced Computer Studies, College Park, MD, USA, 1990.
- [25] ———, *Skip lists: A probabilistic alternative to balanced trees*, Commun. ACM **33** (1990), no. 6, 668–676.

Benjamin A. Burton
 School of Mathematics and Physics, The University of Queensland
 Brisbane QLD 4072, Australia
 (bab@maths.uq.edu.au)

Appendix

In the introduction we claim there are at least 6.44×10^{12} connected combinatorial triangulations of size $n = 9$, up to isomorphism. Here we give the arguments to support this claim.

We begin by placing a lower bound on the number of *labelled* connected combinatorial triangulations. To ensure that each triangulation is connected, we insist that the first face of tetrahedron k is glued to some face chosen from tetrahedra $1, \dots, k-1$, for all $k > 1$. Of course there are many labelled connected triangulations that do not satisfy this constraint, but since we are computing a lower bound this does not matter.

We initially choose gluings for the first face of each tetrahedron $2, 3, \dots, n$ in order. For the first face of tetrahedron k there are $2k$ choices for a partner face—these are the $4(k-1)$ faces of tetrahedra $1, \dots, k-1$ minus the $2(k-2)$ faces already glued—as well as six choices of rotation or reflection. This gives a total of $4 \times 6 \times \dots \times (2n-2) \times 2n \times 6^{n-1}$ possibilities. From here there are $(2n+1) \times (2n-1) \times \dots \times 3 \times 1 \times 6^{n+1}$ ways of gluing together the remaining $2n+2$ faces in pairs, giving a lower bound of at least $(2n+1)! \times 6^{2n}/2$ labelled connected combinatorial triangulations of size n .

We finish by factoring out isomorphisms. Each isomorphism class has size at most $n! \times 4!^n$ (all possible relabellings of tetrahedra and their vertices), and so the total number of connected combinatorial triangulations of size n *up to isomorphism* is at least

$$\frac{(2n+1)! \times 6^{2n}}{2 \times n! \times 4!^n}.$$

For $n = 9$ this evaluates to approximately 6.4435×10^{12} .